

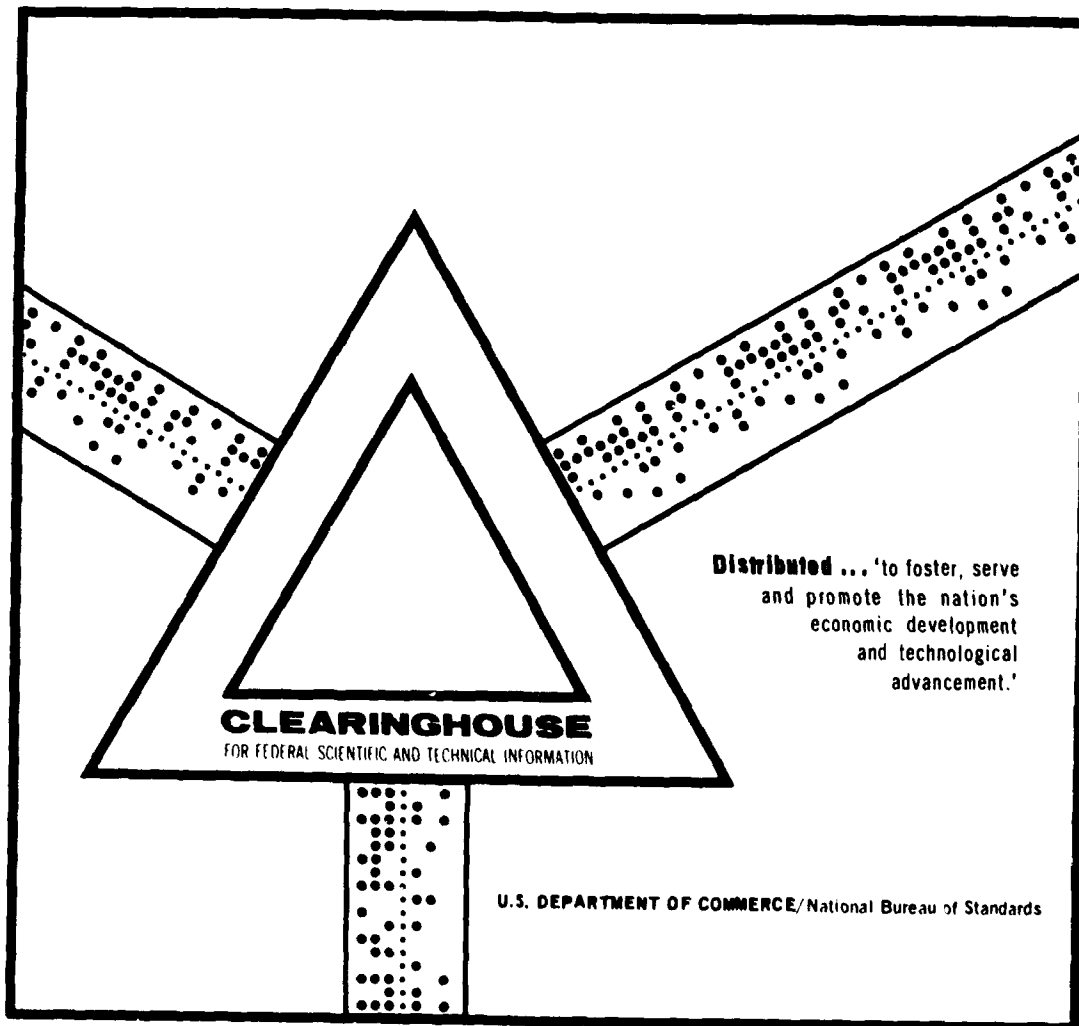
AD 700 763

RESEARCH IN THE CONTROL OF COMPLEX SYSTEMS

Alvin W. Drake

Massachusetts Institute of Technology
Cambridge, Massachusetts

31 May 1967



This document has been approved for public release and sale.

AD700763

FINAL REPORT

"Research in the Control of Complex Systems"

Office of Naval Research

Contract Nonr - 1841(87)

NR 042-230


M.I.T. DSR 9153 and 79153

Report Period: June 1, 1962 - May 31, 1967


Operations Research Center
Massachusetts Institute of Technology
Cambridge, Massachusetts 02139

Author of Report: Alvin W. Drake

Reproduced by the
CLEARINGHOUSE
for Federal Scientific & Technical
Information Springfield, Vt. 01104



B



10

FINAL REPORT TO ONR, Contract Nonr-1841(87)

This report concerns research completed on "Research in the Control of Complex Systems," supported by the Office of Naval Research from June 1, 1962 through May 31, 1967, at the M.I.T. Operations Research Center under Contract Nonr-1841(87), NR 042-230. The contract funds were used to support Scientific Personnel (see Section A below) whose research (Section D) was published (Section B) in the Operations Research Center's technical report series. The contract was also used to support graduate research published as theses (Section C).

This document is a final (summary) report. Ronald A. Howard was the Principal Investigator from the inception of the contract, June 1, 1962 through May 31, 1965. Philip M. Morse was Principal Investigator June 1, 1965 through May 31, 1966. Philip M. Morse and Alvin W. Drake were the Principal Investigators June 1, 1966 through May 31, 1967.

The Progress Reports on this contract may also be of interest.

SECTION A

SCIENTIFIC PERSONNEL (supported entirely or in part)

Prof. Ernest F. Bisbee

Sept. 16, 1962 - June 15, 1963

Prof. Ronald A. Howard

Sept. 16, 1962 - June 15, 1963

July 1, 1963 - August 31, 1963

Sept. 16, 1963 - January 31, 1964

July 16, 1964 - July 30, 1964

Mr. J. David R. Kramer, Jr.

Sept. 16, 1962 - June 15, 1963

Feb. 1, 1964 - June 15, 1964

Prof. Alvin W. Drake

Sept. 16, 1964 - June 30, 1965

Prof. Philip M. Morse

Sept. 16, 1966 - June 15, 1967

Edward A. Silver

July 1, 1962 - August 31, 1962

Sept. 16, 1962 - June 15, 1963

July 1, 1963 - Sept. 15, 1963

Ronald C. Rosenberg

Sept. 16, 1962 - January 31, 1963

SECTION A. Scientific Personnel (continued)

Stephen M. Pölldöck
Sept. 16, 1962 - June 15, 1963

Kendrick B. Melrose
April 1, 1963 - June 15, 1963

Kenneth Scott
Feb. 1, 1963 - June 30, 1963

N. R. Patel
June 16, 1963 - Sept. 15, 1963

Romulo H. Gonzalez
Sept. 16, 1963 - Sept. 15, 1964

Ralph L. Miller
Sept. 16, 1963 - June 15, 1964
Sept. 16, 1964 - June 15, 1965

John U. Beusch
Sept. 16, 1964 - June 15, 1965

Earl D. Brown
Feb. 1, 1965 - June 15, 1965

John M. Cozzolino, Jr.
Sept. 16, 1962 - June 15, 1963
Sept. 16, 1963 - June 15, 1964
Sept. 16, 1965 - Sept. 15, 1966

E. Gerald Hurst, Jr.
Sept. 16, 1965 - Sept. 15, 1966

SECTION B

TECHNICAL REPORTS (supported entirely or in part):

Silver, E. A., "The Transient Solutions for 3-State Discrete Time Markov Processes," Technical Note No. 1, M.I.T. Operations Research Center, July 1963.

Silver, E. A., "Markovian Decision Processes with Uncertain Transition Probabilities or Rewards," Technical Report No. 1, M.I.T. Operations Research Center, August, 1963.

Howard, Ronald A., "Semi-Markovian Control Systems, " Technical Report No. 3, M.I.T. Operations Research Center, December, 1963,

Rothkopf, Michael H., "Scheduling Independent Tasks on One or More Processors, " Technical Report No. 2, M.I.T. Operations Research Center, January, 1964.

Silver, Edward A., "The Use of the Hypergeometric Function as Part of Bayesian Estimation in a Two State Markov Process, " Technical Report No. 2, M.I.T. Operations Research Center, January, 1964.

Kramer, J.D. R., Jr., "Partially Observable Markov Processes, " Technical Report No. 11, M.I.T. Operations Research Center, March, 1965.

Howard, Ronald A., "Dynamic Inference, " Technical Report No. 10, M.I.T. Operations Research Center, December, 1964.

Cozzolino, J. M., R. H. Gonzalez-Zubieta, and R. L. Miller, "Markovian Decision Processes with Uncertain Transition Probabilities, " Technical Report No. 11, M.I.T. Operations Research Center, March, 1965.

Beusch, John U., "Dynamic Behavior and Control of Communications Networks, " Technical Report No. 13, M.I.T. Operations Research Center, June, 1965.

Schweitzer, P. J., "Perturbation Theory and Markovian Decision Processes, " Technical Report No. 15, M.I.T. Operations Research Center, June, 1965.

Miller, R. L., "A Model for Traffic Flows on a Two-Lane Two-Way Rural Highway, " M.I.T. Operations Research Center Technical Report No. 20, June, 1966.

Cozzolino, J. M., Jr., "The Optimal Burn-in Testing of Repairable Equipment, " M.I.T. Operations Research Center, October, 1966.

Hurst, E. G., Jr., "A Class of Models for Adaptive Experimentation and Control, Technical Report No. 30, M.I.T. Operations Research Center, July, 1967.

Brown, E. D., "Some Mathematical Models of Inspection Along a Production Line, " Technical Report No. 36, M.I.T. Operations Research Center, May, 1968.

SECTION C

THESES PUBLISHED (supported entirely or in part):

June 1962

Beusch, John U., "Dynamic Behavior of Linear, Coupled, Chemical Processes," S. M. Thesis, Dept. of Electrical Engineering, supervised by Howard.

Rothkopf, Michael H., "Priority Assignment in Queues with Non-Linear Waiting Costs," S.M. Thesis, Sloan School of Management, supervised by Greenberger.

February 1963

Scott, Kenneth R., "Optimal Multi-Region Discrete Search Under Linear Search Cost," S.M. and E.E. Thesis, Dept. of Electrical Engineering, supervised by Galliher.

September 1963

Silver, Edward A., "Markovian Decision Processes with Uncertain Transition Probabilities," Sc.D. Thesis, Dept. of Civil Engineering, supervised by Howard.

February 1964

Rothkopf, Michael H., "Scheduling Independent Tasks on One or More Processors," Ph.D. Thesis, Sloan School of Management, supervised by Greenberger.

June 1964

Cozzolino, John M., Jr., "Optimum Policies Under Uncertainty," S.M. Thesis, Sloan School of Management, supervised by Howard.

Kramer, J. David R., Jr., "Partially Observable Markov Processes," Ph.D. Thesis, Dept. of Electrical Engineering, supervised by Howard.

Gonzalez-Zubieta, Romulo H., "On Some Aspects of Integer Linear Programming," Ph.D. Thesis, Sloan School of Management, supervised by Little.

Hurst, E. Gerald, Jr., "A Dynamic Programming Formulation for Disposable and Repairable Inventory," S.M. Thesis, Sloan School of Management, supervised by Bisbee.

Section C, Theses Published (continued)

Schweitzer, Paul J., "Perturbation Theory and Markovian Decision Processes, Sc.D. Thesis, Dept. of Physics, supervised by Morse.

June 1966

Miller, Ralph L., "Model for Traffic Flows on a Two-Lane Two-Way Rural Highway," Ph.D. Thesis, Dept. of Mathematics, supervised by Morse.

September 1966

Cozzolino, John M., Jr., "Optimal Burn-in Testing of Repairable Equipment," Ph.D. Thesis, Sloan School of Management, supervised by Little.

February 1968

Brown, Earl D., "Some Mathematical Models of Inspection Along a Production Line," Ph.D. Thesis, Dept. of Mathematics, supervised by White.

SECTION D

SUMMARIES OF RESEARCH FINDINGS (supported entirely or in part):

"Markovian Decision Processes with Uncertain Transition Probabilities or Rewards," by Edward A. Silver, Technical Report No. 1, M.I.T. Operations Research Center, August, 1963 (supervised by Howard).

In most Markov process studies to date it has been assumed that both the transition probabilities and rewards are known exactly. The primary purpose of this thesis is to study the effects of relaxing these assumptions to allow more realistic models of real world situations. The Bayesian approach used leads to statistical decision frameworks for Markov processes.

The first section is concerned with situations where the transition probabilities are not known exactly.

SECTION D, Summaries of Research Findings (continued)

One approach used incorporates the concept of multi-matrix Markov processes, processes where it is assumed that one of several known transition matrices is being utilized, but we only have a probability vector on the various matrices rather than knowing exactly which one is governing the process. An explanation is given of the Bayes modification of the probability vector when some transitions are observed. Next, we determine various quantities of interest, such as mean recurrence times. Finally, a discussion is presented of decision making involving multi-matrix Markov processes.

The second approach assumes more directly that the transition probabilities themselves are random variables. It is shown that the multidimensional Beta distribution is a most convenient distribution (for Bayes calculations) to place over the probabilities of a single row of the transition matrix. Several important properties of the distribution are displayed. Then a method is suggested for determining the multidimensional Beta prior distributions to use for any particular Markov process. Next, we deal with the effects on various quantities of interest of having such distributions over the transition probabilities. For 2-state processes, several analytic results are derived. Despite analytic complexities, some interesting expressions are developed for N-state cases.

It is shown that for decision purposes the expected values of the steady state probabilities are important quantities. For a special 2-state situation, use of the hypergeometric function (previously utilized in the solution of certain physics problems) permits evaluation of these expected values. Their determination for 3 or more states requires the use of simulation. Fortunately, a simple approximation technique is shown to generally give accurate estimates of the desired quantities. An entire chapter is devoted to statistical decisions in Markov processes when the transition probabilities are multidimensional Beta distributed rather than being exactly known. The main problem considered is one where we have the option of buying observations of a Markov process so as to improve our knowledge of the unknown transition probabilities before deciding whether or not to utilize the process.

In the second section of the study, we assume that the transition probabilities are exactly known, but now the rewards are random variables. First, we display the Bayes modification of two convenient distributions to use for the rewards. Next, the expected rewards in various time periods are determined. Finally, an explanation is presented of how to utilize these expected rewards in making statistical decisions concerning Markov processes whose rewards are not known exactly.

SECTION D, Summaries of Research Findings (continued)

"The Transient Solutions for 3-State Discrete Time Markov Processes, " by Edward A. Silver, Technical Note No. 1, M.I.T. Operations Research Center, July, 1963.

Through the combined use of flow graph analysis, geometric transform techniques, and the properties of quadratic functions, we have obtained closed-form expressions for the transient solution of any time-discrete 3-state Markov process. The results are significantly more complex than in the 2-state case. For evaluating the actual multi-step transition probabilities there are five different cases to be considered. The two critical quantities which decide which case is to be used are M and q which, in turn, are involved quadratic functions of the six transition probabilities. Furthermore, there are thirteen different possible general transient behaviors.

Although the calculations are involved, there is one important time-saving factor which should be mentioned. All functions of n are independent of the particular (i, j) pair; hence, need only be calculated once to obtain all nine elements of $\Phi(n)$.

"Semi-Markovian Control Systems, " by Ronald A. Howard, Technical Report No. 3, M.I.T. Operations Research Center, December, 1963.

Decision models of the type we have discussed are interesting mathematical developments, but they are useful in practical problems only if we can supply them with necessary data. Consequently, we are investigating schemes for deriving this data from both experience and experiment. We are also investigating how to solve the decision process when the data are uncertain. Although the practicality of these methods will be considerably enhanced by the completion of this research, the model as it stands continues to have important application in several areas of management systems.

"Scheduling Independent Tasks on One or More Processors, " by Michael H. Rothkopf, Technical Report No. 2, M.I.T. Operations Research Center, January, 1964 (supervised by Greenberger).

Although much attention has been given to problems relating to scheduling dependent tasks, relatively little attention has been given to the simpler problem of scheduling independent tasks. This thesis considers the problem of scheduling m independent, immediately available tasks on n processors. Each task, i , has a service time, T_i , and a waiting cost rate, $c_i(t)$, that is a function of time. There are no feasibility restrictions on the order in which the tasks are to be processed. The problem is related to the job shop scheduling problem and the problem of optimally assigning priorities in a queuing system.

SECTION D, Summaries of Research Findings (continued)

It is shown that when $c_i(t) = c_i e^{-rt}$ (i.e., linear waiting costs with continuous discounting at rate r) and there is a single processor waiting, costs are minimized by sequencing the tasks in decreasing $rc_i e^{-rT_i} / (1 - e^{-rT_i})$ order.

A dynamic programming algorithm has been developed for a wide class of multiprocessor scheduling problems. Linear waiting costs (with or without discounting), arbitrary processor use costs, and certain other costs are considered. In certain cases, the service time of a task may be a function of the processor that performs it.

Several algorithms are presented for scheduling tasks with absolute deadlines on one or more processors. For certain classes of problems, the unprofitability of splitting tasks in time and/or between processors is proved. Scheduling when service times and cost rates are known only stochastically is discussed. Many results for non-stochastic scheduling are extended. A simple technique is presented for finding "near optimal" sequences for a single processor and tasks with arbitrary waiting cost functions.

In appendices, dynamic programming algorithms are presented for a scheduling problem with interruptions and for a generalized traveling salesman problem.

"The Use of the Hypergeometric Function as Part of Bayesian Estimation in a Two State Markov Process," by Edward A. Silver, Technical Note No. 2, M.I.T. Operations Research Center, January 1964.

This paper deals with a special case of Bayesian estimation in Markov processes. We consider a two-state process where one transition probability is known exactly, while the other is assumed to be Beta distributed. Under these conditions the expected values of the steady state probabilities are obtained through the use of the hypergeometric function, a mathematical function heretofore encountered only in an entirely different area of applied mathematics. Knowing the expected values of the steady state probabilities enables us to place the process considered into a statistical decision framework.

"Partially Observable Markov Processes," By J. David R. Kramer, Jr., Technical Report No. 4, M.I.T. Operations Research Center, April, 1964, (supervised by Howard).

SECTION D, Summaries of Research Findings (continued)

A partially observable Markov process is a model of a discrete time dynamic system which takes into account the effects of imperfect observations and of random system behavior.

The model consists of an underlying Markov process with state vector $X(n)$. Direct observations of $X(n)$ are not possible, but a vector $Z(n)$ is observed. The observation $Z(n)$ is related to the state $X(n)$ by a known probability density function $\gamma(Z(n)|X(n))$. This model is useful in the analysis of a very large class of sequential decision problems.

It is shown that a partially observable Markov process is conveniently analyzed by the introduction of the probability density function $f(X(n)|O(n))$, where $O(n)$ represents the aggregate of observations made on the system up to and including time n . This density function is shown to have certain characteristic iterative properties and is referred to as the statistical state of the system. In particular it is shown that $f(X(n+1)|O(n+1))$ may be calculated from a knowledge of $f(X(n)|O(n))$ as a function of $X(n)$ and the observation $Z(n+1)$.

In any practical situation it is necessary to replace the density function $f(X(n)|O(n))$ by a finite set of parameters. Several methods of carrying out this parameterization are discussed.

The application of the theory of partially observable Markov processes to the problems of estimation, prediction, and smoothing is straightforward. When a general terminal control problem is considered, however, the notion of "minimum expected cost" turns out to be ambiguous. The concepts of a priori and a posteriori controls are distinct. Specific numerical results from a computer program written to determine the optimal control are given.

"Dynamic Inference," by Ronald A. Howard, Technical Report No. 10, M.I.T Operations Research Center, December, 1964.

One of the problems that continually arise in analyzing systems is that of modeling situations where the underlying statistical parameters of the process may change from time to time. Information on the changes may be revealed only through certain observable variables that are statistically related to the parameters. We must infer from the pattern of observable variables the probability distributions of process parameters and then use these distributions to predict the future course of the process; we call this process "Dynamic Inference."

For a military illustration of this problem, let us consider the case of detecting submarines using a fixed sonar station. If a submarine is

SECTION D, Summaries of Research Findings (continued)

present in a station's detection area, each sonar return pulse energy will be selected from one probability distribution, while if a submarine is not present the pulse energy will be selected from another probability distribution. The time between arrivals of submarines to the station's area is one random variable; the length of time they stay after they get there is another. The observables in the process are the sequence of sonar returns, which we shall call the observable pattern. The problem is to assign a probability to a submarine's being present in the detection area, given an observable pattern. This probability is naturally modified as successive sonar pulses are received. On the basis of the probability and associated system operating costs, the decision maker will then decide at any point whether to send an investigating aircraft or destroyer for further confirmation of the submarine's presence.

"Markovian Decision Processes with Uncertain Transition Probabilities," by John M. Cozzolino, Romulo H. Gonzalez-Zubieta, and Ralph L. Miller, Technical Report No. 11, M.I.T. Operations Research Center, March, 1965.

The theory for finding optimal policies for Markov processes with transition rewards and many alternatives in each state, when the transition probabilities are given, has already been developed. But in most practical applications, these transition probabilities are not known exactly--one has only some prior knowledge about them.

This problem of uncertain transition probabilities was first treated by Dr. F. A. Silver in Interim Technical Report No. 1 of the M.I.T. Operations Research Center. Section I of the present report extends several of Silver's results. We propose a dynamic programming formulation for the problem of choosing an optimal operating strategy, and we carry out the solution for a special two-state example. However, it is found that the solution of non-trivial problems of any higher dimension is impractical.

Section II is concerned with experimental and heuristic approaches to the problem, and relies upon simulation rather than upon analysis. We investigate certain statistics of the process when the unknown transition probabilities are governed by a multidimensional beta prior (a convenient form for Bayes modification). We find that the process with known probabilities which are equal to the mean values of the unknown probabilities, provides us with a remarkably good picture of the unknown process. The hypothesis is stated and investigated that this process of expected values is adequate for decision purposes, and that determining decisions from it is feasible as well as useful.

Finally, some alternative approaches are suggested for cases which might not be handled by the expected-values technique.

SECTION D, Summaries of Research Findings (continued)

"Dynamic Behavior and Control of Communications Networks," by John U. Beusch, Technical Report No. 13, M.I.T. Operations Research Center, June, 1965 (supervised by Howard).

An arbitrary store-and-forward network where messages without priorities can enter the network at any node and be destined for any other node is considered. The problem of routing messages is considered to be a feedback control problem in which information about the number of bits waiting at various places in the network is fed back to points where decisions are made.

Mathematical models which represent system behavior are developed. The model of a single channel is a new Markov process, ~~queuing model~~ with random arrivals and deterministic service times. Simple models which predict future queue lengths are developed. A way of reducing a set of series-parallel channels to a single equivalent channel is presented.

A method of determining message routing decision rules which are based on what is known at a decision point is developed. A method of determining whether and how often to obtain information about another part of the network is developed.

Many of the methods apply only to series-parallel networks.

Several examples to illustrate the methods and a computer program for analyzing networks are presented. Results of simulations to compare various routing techniques with the feedback techniques are presented. In all cases feedback improves performance and when average arrival rates of messages change it improves performance drastically.

"Perturbation Theory and Markovian Decision Processes," by Paul J. Schweitzer, Technical Report No. 15, M.I.T. Operations Research Center, June, 1965 (supervised by Morse).

The Howard-Jewell algorithm for programming over a Markov-renewal process is analyzed in terms of a perturbation theory formalism which describes how the stationary distribution changes when the transition probabilities change. The policy improvement technique is derived from this viewpoint. The relative values may be interpreted as partial derivatives of the gain rate with respect to policy.

The value equations are shown to be solvable, with the relative values unique up to one additive constant, if and only if the underlying Markov chain is irreducible. The policy iteration algorithm is shown not to cycle, thus guaranteeing convergence.

SECTION D, Summaries of Research Findings (continued)

A discussion of the existence, uniqueness, and characterization of the solution to the functional equation of dynamic programming is given. Emphasis is placed upon the value-maximization of transient states.

The fundamental matrix is developed as a useful tool for doing perturbation theory, describing first-passage properties of semi-Markov processes, and for dealing with semi-Markov processes with rewards.

"A Model for Traffic Flows on a Two-Lane Two-Way Rural Highway," by Ralph L. Miller, Technical Report No. 20, M.I.T. Operations Research Center, June, 1966 (supervised by Morse).

One is given a distribution of desired, or free, speeds (the speed a driver would maintain if he were unimpeded by traffic). The total density of traffic in both lanes is also known. The problem is to predict the steady state values of the following basic quantities: (1) the distribution of forced speed (the speed a driver is actually able to maintain in traffic) as a function of the desired speed; and (2) queue composition (the number of cars of each desired speed in a queue) as a function of the speed of the queue leader. Once these basic quantities are computed, it is easy to deduce: (3) mean delays experienced by drivers of each free speed; (4) the distribution of queue speeds; and (5) the total flow of traffic. Other statistics of interest also follow readily. In fact, it is asserted, the basic quantities are so basic that they tell us all we could ever want to know about the road in steady state.

This paper presents a method for completely determining both basic quantities, subject to certain assumptions which are made explicit. No attention is given to the process by which the steady state is reached. Moreover, since the model to be used is stochastic, it applies only to relatively low traffic densities. For higher densities, a dynamic car-following model is more appropriate. The limits in which this model is valid is discussed.

Computational questions are also considered and several sample calculations are made. The results are compared with empirical data to measure the validity of the model. Finally, various applications of the model are explored, including highway safety, setting legal speed limits, and automotive design criteria.

The paper concludes with a discussion of the limitations of the model, and suggestions are made for future research.

SECTION D, Summaries of Research Findings (continued)

"The Optimal Burn-In Testing of Repairable Equipment," by John M. Cozzolino, Technical Report No. 23, M.I.T. Operations Research Center, October, 1966 (supervised by Little).

The infant mortality effect observed in the statistical treatment of reliability denotes a decreasing with time of the conditional probability of failure of a device which exhibits it. This widely present effect may be utilized to improve the reliability by means of burn-in testing which seeks to discriminate between high-and-low-quality units by accumulating operating experience upon all units.

Burn-in testing is applicable to both unrepairable and repairable devices. The latter case is more difficult since the state of a failed and repaired unit usually depends upon the entire past history of the unit and upon the nature of the repair process.

A general burn-in test problem for repairable devices is formulated based upon the explicit modeling of the repair process. The case of unrepairable devices is treated as a special case. Considering a particular conjugate form of the repair rule, the burn-in test optimization is formulated as a sequential decision problem and the solution is discussed in terms of dynamic programming.

An adaptive reformulation which allows a parameter of the failure process to be considered unknown is also given.

Specific models of decreasing failure rate processes, based upon the population heterogeneity cause of decreasing failure rate, are given.

"A Class of Models for Adaptive Experimentation and Control," by E. Gerald Hurst, Jr., Technical Report No. 30, M.I.T. Operations Research Center, July, 1967 (supervised by Little).

Solutions are presented for several control problems with discretely dynamic, stochastic, partially observable states in which the amount of experimentation at each stage constitutes an important control decision.

Bayesian autoregressive time-series models are given, both in general and assuming normal density function for the change process, data generator, and statistical description of the state. These include first-order and k^{th} -order autoregressive models both with and without unknown parameters.

A general dynamic programming formulation for control of the known first-order process is obtained; it is specialized to the case with

SECTION D, Summaries of Research Findings (continued)

quadratic cost of error and proportional cost of experimentation. The optimal experimental policy at every stage is found to consist of a single critical number; if the precision of current information is below this level, experimentation is performed. It is noted that this structure corresponds to the optimal ordering policy for inventories with proportional ordering cost and convex operating cost. The form of and numerical values for the steady state policy are presented.

A set-up cost of experimentation is introduced, and a two-level experimental policy analogous to the (s, S) policy in inventory theory is obtained. The assumption of completely known process parameters is relaxed by allowing uncertainty in the precision of change. A three-variable dynamic programming formulation which simultaneously infers the current level of the physical state and the unknown precision is solved for its optimal experimental policy, which is described by two critical numbers. A simple approximately optimal policy in terms of the earlier numerical results is proposed.

"Some Mathematical Models of Inspection Along a Production Line," by Earl D. Brown, Technical Report No. 36, M.I.T. Operations Research Center, May, 1968 (supervised by White).

A production line is taken to be an ordered series of manufacturing operations separated by potential inspection stations. The study concerns two primary aspects of inspection along a production line through which item goods are processed. First, when the system characteristics which affect final quality are not known with certainty, the problem of allocating costly inspection throughout the line is considered. Two models are presented, analyzed, and combined. Dynamic programming formulations are closely analyzed and much attention is given to securing bounds on their operation and on the optimal solution. A case is demonstrated where two such formulations may be combined to the advantage of each. Several relationships are established between the optimal solution for the complete line and that for certain truncated lines, and computer-run examples are given. The analysis is extended to cases of arbitrary inspection cost patterns and setup costs. A formulation occurs resembling the dynamic inventory problem and a rigorous proof is given of Scarf's result concerning the structure of the optimal solution.

Secondly, when the system characteristics are known with certainty, there is the problem of installing a combined inspection and disposition plan in order to "control" the final quality. Two models are described and for each the properties of the optimal solutions are discussed. For a "long-term-cost" model, an "all or nothing" property is proved for the optimal inspection as well as for the optimal disposition plans, but the question is

SECTION D, Summaries of Research Findings (continued)

reopened for the "long-run-quality" model. Various classes of quality control procedures and costs are considered.

Finally, several possibilities for further research are inferred from both aspects of the work.